Cascaded FICON: A Brief Tutorial

By Steve Guendert and Tedd Cesiano

INTRODUCTION

FICON, like most technological advancements, evolved from the limitations of its predecessor. ESCON (IBM's Enterprise System Connection) is a very successful storage network protocol for mainframe systems and has been considered the father of the modern SAN. Since its introduction in 1999, FICON has seen quite an evolution in its brief history. We started out with FICON bridge mode (FCV), and in 5 short years have gone to single director Native FICON (FC) implementations, to configurations that intermix FICON and open systems Fibre Channel Protocol (FCP). The most recent addition to the capabilities of FICON is the ability to create cascaded fabrics of FICON directors and switches.

Cascaded FICON allows the end user to have a FICON channel, or FICON CTC to connect an IBM zSeries server to another zSeries server or peripheral device such as disk, tape library, or printer via two FICON directors or switches between the connected devices and/or servers. This allows for tremendous flexibility and fabric cost savings in the FICON architecture, better utilization of storage resources, and higher data availability in the enterprise. It also allows for more robust disaster recovery and business continuity.

This article will discuss what cascaded FICON is, how it works, and what goes into the planning, design and implementation of a cascaded FICON environment.

14 YEARS OF EVOLUTION

The ESCON channel architecture was introduced to the world in 1990 as the way to address the limitations of parallel (bus and tag) architectures. As such, ESCON provided noticeable, measurable improvements in distance capabilities, switching topologies, and most importantly, response time and service time performance. By the end of the 1990s, ESCON's strengths over parallel channels had become its weaknesses. FICON evolved in the late 1990s to address the technical limitations of ESCON in bandwidth, distances, and channel/device addressing.

Initially, the FICON (FC-SB-2) architecture did not allow the connection of multiple FICON directors. Of course, neither does ESCON except when static connections of "chained" ESCON directors were used to extend ESCON distances. Both ESCON and FICON defined a single byte for the link address, the link address being the port attached to "this" director. As of 31 January 2003 this changed. Now, it is possible to have two-director configurations, as well as separate geographic sites. This is done by adding the domain field of the fibre channel destination ID to the link address. This is done to specify the exiting director and the link address on that director.

WHAT IS CASCADED FICON?

Cascaded FICON refers to an implementation of FICON that involves one or more FICON channel paths to be defined over 2 FICON directors that are connected to each other using an Inter-Switch Link (ISL). The processor interface is connected to one director, while the storage interface is connected to the other. This configuration is supported for both disk and tape, with multiple processors, disk subsystems and tape subsystems sharing the ISLs between the directors. Multiple ISLs between the directors are also supported.

There are some hardware and software requirements specific to cascaded FICON:

1. The FICON directors themselves must be from the same vendor (i.e. both from Brocade, Cisco, CNT, or McDATA)
2. The mainframes must be zSeries machines: z800, 890, 900, 990. Cascaded FICON requires 64-bit architecture to support the two-byte addressing scheme. Cascaded FICON is therefore not supported on 9672 G5/G6 mainframes.
3. z/OS version 1.4 or greater, and/or z/OS version 1.3 with required PTFs/MCLs to support 2-byte link addressing (DRV3g and MCL (J11206) or later).
4. The high integrity fabric feature for the FICON director must be installed on all directors involved in the cascaded architecture.

WHAT ARE THE BENEFITS?

Cascaded FICON delivers many of the same benefits of open systems SANs to the mainframe space. Cascaded FICON allows for simpler infrastructure management, lowered infrastructure cost of ownership, and higher data availability. This higher data availability is important in delivering a more robust enterprise disaster recovery strategy. The benefits are further realized when the ISLs connect directors in two or more locations and/or are extended over long distances. Figure 1 shows a non-cascaded two-site environment.

In the configuration in Figure 1, all hosts have access to all of the disk and tape subsystems at both locations. The host channels at one location are extended to the FICON directors at the other location to allow for cross-site storage access. If each line represents two FICON channels, then the above configuration would need a total of sixteen (16) extended links. These links would only be utilized to the extent that the host has activity to the remote devices.

The most obvious benefit when comparing the above configuration with one that is cascaded is the reduction in the number of links across the WAN. The following diagram shows a cascaded, two-site FICON environment.

In this configuration, if each line represents two channels, only four (4) extended links are required. Since FICON is a packet-switched environment (as opposed to the circuit-switched ESCON environment), multiple devices can share the ISLs, and multiple I/Os can be processed across the ISLs at the same time. This allows for the reduction of links between sites and allows for more efficient utilization of the links in place. In addition, ISLs can be added as the environment grows and traffic patterns dictate.

This is the key way in which a cascaded FICON implementation can reduce the cost of the enterprise architecture. As can be seen in the above diagrams, the cabling schema both intersite and intrasite has been simplified. Fewer intrasite cables translate into lower cabling hardware and management costs. It also lowers the number of FICON adapters, director ports, and host channel card ports required, thus lowering the connectivity cost for your mainframes and storage devices as well. As you can see in FIGURE 2, the sharing of links between the two sites reduces the number of physical channels going between sites, again, lowering the cost by consolidating channels, and the number of director ports. Incidentally, the faster the channel speeds between sites, the better the intersite cost savings from this consolidation will be. So, as 4 Gbps FICON, and 10 Gbps FICON become available, the more attractive this arrangement becomes.

Another benefit to this approach, especially over long distances, is that the FICON director typically has many more buffer-credits per port than do the processor and the disk/tape subsystem cards. More buffer-credits allow for a link to be extended farther distances without significantly impacting response times to the host. The following table illustrates how buffer credits affect the allowable distance for each link speed, using this formula:

Distance (KM) = (Buffers - 1) * (2 / link speed in Gb)

Notice that with high-speed links and low numbers of buffer-credits, the channel cannot be extended very far.

What does this really mean as a benefit? In a nutshell it means that cascaded FICON allows for even higher availability, disaster recovery, and business continuity (HA/DR/BC).

A FEW WORDS ON HA/DR/BC

The greater bandwidth of and distance capabilities FICON has over ESCON are starting to make it an essential and cost effective component in HA/DR/BC solutions. Since September 11, 2001 more and more companies are insourcing DR, and those that are doing so are building the mainframe piece of their new DR/BC datacenters using FICON, rather than ESCON. Until IBM announced FICON cascading as being GA, the FICON architecture has been limited to a single domain due to the single byte addressing limitations inherited from ESCON. FICON cascading allows the end user to have a greater maximum distance between sites: up to an unrepeated distance of 36 km at 2 Gb/sec bandwidth.

September 11 drove home how critical it is for an enterprise to be prepared for disaster. Even more so for large enterprise mainframe customers. A complete paradigm shift has occurred since 9/11/01 when we discuss DR/BC. Disaster recovery is no longer limited to problems such as fires or a small flood. Companies now need to consider and plan for the possibility of the destruction of their entire data center, and possibly the people that work in it. A great many articles, books, and other publications have discussed the IT "lessons learned" from September 11, 2001:

FIGURE 1: TWO-SITE, NON-CASCADED FICON ENVIRONMENT

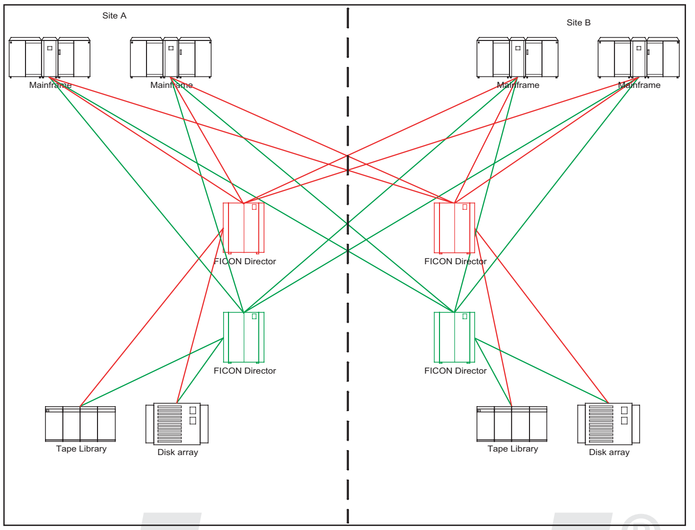

FIGURE 2: TWO-SITE, CASCADED FICON ENVIRONMENT

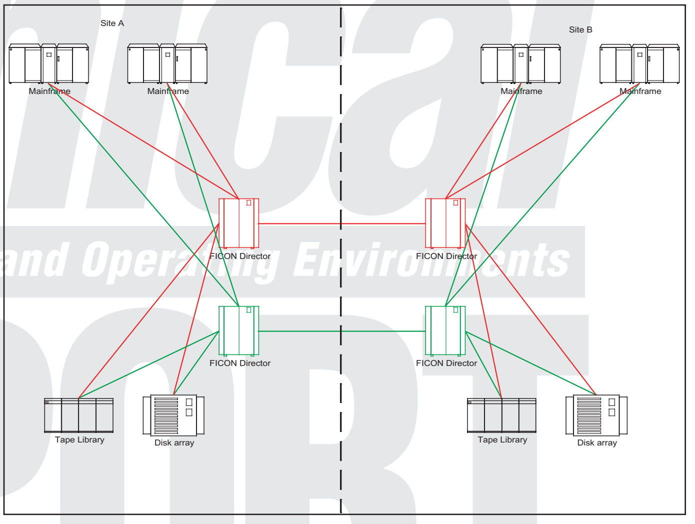

1. To maintain business continuity it is absolutely critical to maintain geographical separation of facilities and resources. Any resource your enterprise has that cannot be replaced from external sources within your recovery time objective (RTO) should be available within the enterprise. It is also preferable to have these resources in multiple locations. We're talking about buildings, hardware, software, data, and staff. Cascaded FICON gives the geographical separation that post 9/11 business requires; ESCON does not.
2. The most successful DR/BC implementations are often based on as much automation as possible. 9/11/01 proved that key staff and skills may no longer be present after a disaster strikes.
3. Financial, government, military, and other enterprises now have critical Recovery Time Objectives that are measured in seconds

or minutes and not days and hours. For these end users it has become increasingly necessary to implement an in-house (insourced) DR solution. By in-house we mean that the facilities and equipment needed to achieve the HA/DR/BC solution are owned by the enterprise itself. Cascaded FICON allows for considerable cost savings compared with ESCON when insourcing HA/DR/BC.

TECHNICAL DISCUSSION ON FICON CASCADING

The Basics

First, as we stated earlier, cascaded FICON is limited to zSeries mainframes only. Please see the HW/SW requirements outlined earlier in the article. For more details on some of the fibre channel terminology that is in this section, please refer to the key references listed at the end of the article.

Referring to FIGURE 2, you can see that a cascaded FICON director configuration involves at least three fibre channel links. The first link is between the FICON channel card (known as an N_Port) and the FICON director's fibre channel adapter card (which is considered an F_Port). The second link is between the two FICON directors via what are known as E_Ports. The link between E_Ports on the directors is known as an inter-switch link, or ISL. The final link is from the F_Port to a FICON adapter card in the control unit port (N_Port) of the storage device. The physical paths are the actual fibre channel links connected by the FICON directors providing the physical transmission path between a channel and a control unit. Also note that the links between the cascaded FICON directors may be multiple ISLs, both for redundancy and to ensure adequate I/O bandwidth.

Addressing Support

Single byte addressing refers to the link address definition in the Input-Output Configuration Program (IOCP). Two-byte addressing (cascading) allows IOCP to specify link addresses for any number of domains by including the domain address with the link address. This allows the FICON configuration the capability of creating definitions in IOCP that span more than one director.

FIGURE 4 shows that the FC-FS 24 bit FC port address identifier is divided into three fields:

1. Domain
2. Area
3. AL Port

In a cascaded FICON environment, 16 bits of the 24-bit address must be defined for the zSeries server to access a FICON control unit. The FICON directors provide the remaining byte used to make up the full 3-byte FC port address of the CU being accessed. The AL_Port (arbitrated loop) value is not used in FICON and will be set to a constant value. The zSeries "domain" and "area" fields are referred to as the F_Port's "port address" field. It is a 2-byte value, and when defining access to a CU that is attached to this port using the zSeries Hardware Configuration Definition (HCD) or IOCP, the port address is referred to as the link address. FIGURE 5 further illustrates this. FIGURE 6 is an example of a cascaded FICON IOCP gen.

The connections between the two directors are established through the exchange of link parameters (ELP). The directors then pause for a device fabric login (FLOGI), and then, assuming the device is another switch they initiate an ELP exchange. This results in the formation of the ISL connection(s).

FIGURE 3: ALLOWABLE DISTANCES BASED ON CREDITS AND LINK SPEED

Speed	Credits	Distance (km)
1	16	30
2	16	15
4	16	7.5
10	16	3
1	64	126
2	64	63
4	64	31.5
10	64	12.6
1	128	254
2	128	127
4	128	63.5
10	128	25.4

FIGURE 4: FABRIC ADDRESSING SUPPORT (1)

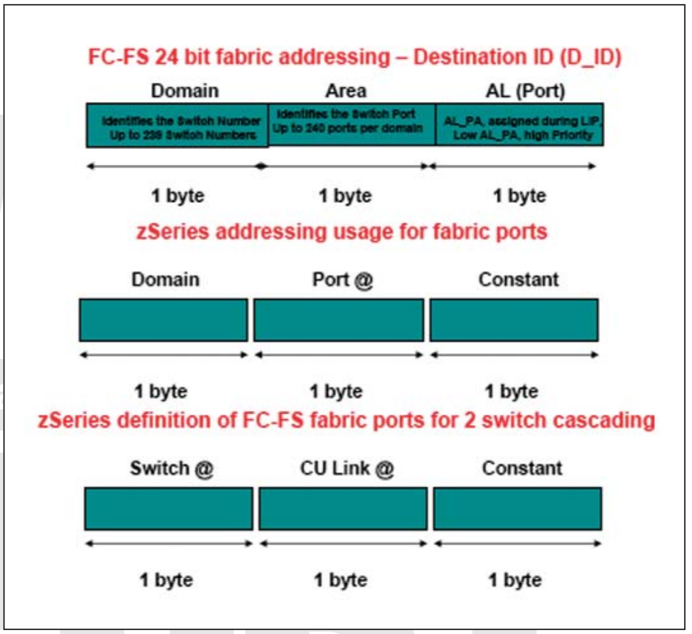

FIGURE 5: FABRIC ADDRESSING SUPPORT (2)

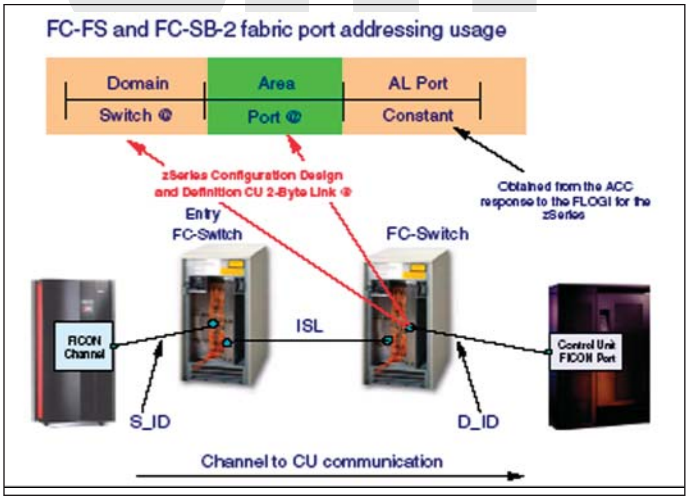

In a cascaded FICON configuration, three additional steps occur beyond the "normal" FICON switched point-to-point communication initialization. A much more detailed discussion of the entire FICON initialization procedure can be found in Chapter 3 of the IBM Redbook: FICON Native Implementation and Reference Guide, pp 23-43. A hyperlink is given in the references. The 3 steps are:

1. If a 2-byte link address is found in the CU macro in IOCDS, a "Query Security Attribute" (QSA) command is sent by the host to check with the fabric controller on the directors if the directors have the high integrity fabric features installed.
2. The director responds to the QSA.
3. If it is an affirmative response, indicating that a high integrity fabric is present (fabric binding and insistent domain ID) the login continues. If not, login stops and the ISLs are treated as invalid (not a good thing).

HIGH INTEGRITY ENTERPRISE FABRICS

Data integrity is paramount in a mainframe environment. End to end data integrity absolutely must be maintained throughout a cascaded FICON director environment. Why? We must ensure that any changes made to the customer's data stream are always detected and that the data is always delivered to the correct end point. What does high integrity fabric architecture and support entail?

1. Support of Insistent Domain IDs. This means that a FICON director will not be allowed to automatically change its address when a duplicate switch address is added to the enterprise fabric. Intentional manual operator action is required to change a FICON director's address. Insistent Domain IDs prohibit the use of dynamic Domain IDs, thereby ensuring that predictable Domain IDs are being enforced within the fabric. It also makes certain that duplicate Domain IDs are not used within the fabric.
2. Fabric Binding. Fabric binding enables companies to allow only FICON directors that are configured to support high-integrity fabrics to be added to the storage/FICON network. The FICON directors that you wish to connect to the fabric must be added to the fabric membership list of the directors already in the fabric. This membership list is composed of the "acceptable" FICON director's WorldWideName (WWN) and Domain ID. Using the Domain ID ensures that there will be no address conflicts, i.e. duplicate domain IDs when the fabrics are merged. The two connected FICON directors then exchange their membership list. This membership list is a Switch Fabric Internal Link Service (SW_ILS) function, which ensures a consistent and unified behavior across all potential fabric access points.

APPLICATIONS FOR CASCADED FICON DIRECTORS

Any application or environment that requires cross-site FICON or fibre channel connectivity for multiple channels can benefit from using cascaded directors. Those that come to mind first are GDPS and array-based replication. Others include remote tape vaulting, where the backup application writes directly to tape drives at the remote location.

FIGURE 6: SAMPLE IOCP GEN FOR CASCADED FICON (Z890/Z990)

```
ID (no change in ID statement for FICON)
RESOURCE PARTITION=((CSS(0),(SYSA,1),(SYSB,2),(SYSC,3)),
      (CSS(1),(SYSD,1),(SYSE,2),(SYSF,3)))

CHPID PATH=(CSS(0),80),SHARED,PART=((CSS(0),(SYSA,SYSB),(=))),TYPE=FC,SWITCH=61,PCHID=160
CHPID PATH=(CSS(0),81),PART=((CSS(0),(SYSB),(=))),TYPE=FC,SWITCH=61,PCHID=1A0
CHPID PATH=(CSS(0),82),PART=((CSS(0),(SYSB),(=))),TYPE=FC,SWITCH=61,PCHID=130
CHPID PATH=(CSS(0,1),83),SHARED,PART=((CSS(0),(SYSC),(=)),(CSS(1),(SYSD),(=))),
      TYPE=FC,SWITCH=61,PCHID=1D0

CNTLUNIT CUNUMBR=8000,PATH=((CSS(0),80,81,82,83),(CSS(1),83)),
      UNITADD=((00,256)),LINK=((CSS(0)6212,6222,6232,6242),(CSS(1),6242)),
      CUADD=0,UNIT=2105
CNTLUNIT CUNUMBR=8100,PATH=((CSS(0),80,81,82,83),(CSS(1),83)),
      UNITADD=((00,256)),LINK=((CSS(0)6212,6222,6232,6242),(CSS(1),6242)),
      CUADD=1,UNIT=2105
        *
        *
CNTLUNIT CUNUMBR=8700,PATH=((CSS(0),80,81,82,83),(CSS(1),83)),
UNITADD=((00,256)),LINK=((CSS(0)6212,6222,6232,6242),(CSS(1),6242)),
      CUADD=7,UNIT=2105

IODEVICE  (no change for FICON)
```

CONCLUSION

The evolution of FICON to support cascading is a clear-cut example of a protocol that was designed to fill ever-changing requirements. FICON is the basis on which mainframe storage networks will be built well into the future. As an upper level layer in the Fibre Channel standard, FICON will continue to be around, most likely evolving into a shared mainframe/open systems storage network.

REFERENCES

- FICON Implementation Guide
  www.redbooks.ibm.com/pubs/pdfs/redbooks/sg246497.pdf
- FICON Native Implementation and Reference Guide
  www.redbooks.ibm.com/pubs/pdfs/redbooks/sg246266.pdf
- Introduction to IBM S/390 FICON
  www.redbooks.ibm.com/pubs/pdfs/redbooks/sg245176.pdf
- Fibre Channel: A Comprehensive Introduction. Robert Kembel.

NaSPA member Steve Guendert is CNT Technical Solutions Group's Storage Principal Consultant and is the Director of their FICON practice. He is a PMI certified Project Management Professional (PMP), and he has an MBA, and a M.S. in Management Information Systems. He is a well regarded FICON consultant and is one of the industry's experts on the subject. He lectures and consults regularly on FICON technology as well as business case justifications for migrating from ESCON to FICON. He is currently a Ph D Candidate and is writing his doctoral dissertation on Enterprise I/O Subsystems.

NaSPA member Tedd Cesiano is a Principal Consultant in CNT's Technical Solutions Group. He has over 15 years of experience in storage management and storage consulting for both mainframe and open systems. He has performed numerous storage, ESCON/FICON, and bandwidth assessments and is the creator of CNT's ESCON/FICON migration assessment deliverable.

